 AIAA 2003-6758

DISTRIBUTED SCHEDULING ARCHITECTURE FOR MULTI-CENTER TIME-BASED METERING

Steven Landry[*], Todd Farley[†], John Foster[‡], Steve Green[§], Ty Hoang[*], Gregory L. Wong[**]
NASA Ames Research Center

ABSTRACT

The Traffic Management Advisor (TMA) is an air traffic control automation system currently in use in seven Air Route Traffic Control Centers (ARTCCs) to enable time based metering to busy airports within their airspace. However, this system is limited to operation within a single ARTCC, within about a 200 nautical mile radius of the airport, and on relatively simple streams of traffic. The need for coordinated metering within a greater (300+ nautical mile) radius of an airport, on streams of traffic with significant branching, and across ARTCC boundaries, has been identified.

Early tests revealed that TMA could not simply be scaled up to handle such a problem. Instead, a loosely coupled hierarchy of schedules, in which constraints from downstream schedules are passed upstream, is required. Such an architecture reduces the reliance on distant projections of arrival times, making schedules robust to changes in sequence and to additions of aircraft (such as aircraft departing inside the system's scheduling horizon). This architecture is also scaleable, easily reconfigurable, and can be networked together. As such, it can be adapted for use in any size or configuration of airspace and with any number of airports delivering restrictions. An implementation of this distributed scheduling architecture is currently undergoing testing in the TMA-Multi Center system. This paper describes the architecture and its motivation.

INTRODUCTION

Time-based metering is a traffic management technique used to stretch out arrival demand during periods when it is expected to exceed the capacity of some National Airspace System (NAS) resource, such as a major airport. As such, metering is an alternative to strategies such as miles-in-trail (MIT) spacing or managed arrival reservoirs. Time-based metering uses automation to assign crossing times for aircraft at points along its route of flight. These "scheduled times of arrival" (STAs) are intended to delay aircraft, providing an efficient and orderly flow of traffic into the constrained resource, while also fully utilizing the available capacity of the resource.

The most widely accepted implementation of time-based metering is the Traffic Management Advisor – Single Center (TMA-SC). TMA-SC was designed and developed by researchers at NASA Ames Research Center in the mid-1990s as part of the Center-Tracon Automation System (CTAS). TMA-SC is in operation at seven Air Route Traffic Control Centers (ARTCCs): Los Angeles, Oakland, Denver, Minneapolis, Fort Worth, Atlanta, and Miami.

Time-based metering, accomplished through the use of TMA-SC, has been effective at increasing capacity and reducing workload at the ARTCCs in which it is deployed. When using TMA-SC, ARTCCs have seen an increase in landing rates of up to 5%. In addition, less holding has been needed, and delays have been shifted further from the TRACON and to higher altitudes. This has resulted in significant fuel savings and has created a smoother flow of aircraft to the airport.[1]

In considering how to expand time-based metering (and its benefits) to more congested airspace, researchers at NASA's Ames Research Center (in collaboration with MITRE's Center for Advanced Automation Systems) examined the feasibility of adapting TMA-SC to the Northeast corridor.

Several major problems were identified.[2] These problems required that a new architecture for time-based metering be developed for airspace such as the Northeast corridor. Researchers at NASA Ames, in concert with engineers at Computer Sciences Corporation, have developed an innovative new architecture for time-based metering to overcome the limitations of the current technology for application to arrival metering

[*] Aerospace Engineer, Member AIAA.
[†] Aerospace Engineer, Senior Member AIAA
[‡] Raytheon ITSS, Deputy Technical Area Manager for CTAS
[§] Manager, En route Systems and Operations, Member AIAA, Associate Fellow
[**] Computer Engineer

within the congested Northeast corridor. The approach makes use of a distributed network of loosely-coupled schedulers to distribute delay over a large area while mitigating errors in predicting trajectories over such distances.

The remainder of this paper will describe this "distributed scheduling" concept. First, some background information will be provided on TMA-SC and the system in which the new time-based metering concept is being tested: Traffic Management Advisor – Multi Center (TMA-MC). Next, a description of the general architecture of the concept will be described. A description of the implementation of the concept within TMA-MC and the airspace in the Northeast corridor is then provided, followed by a discussion of future work and conclusions.

BACKGROUND

TMA-SC scheduling

TMA-SC utilizes centralized scheduling done by a single process, called the dynamic planner (DP). [3] The DP schedules all aircraft landing at an airport, providing an STA for each aircraft at a meter fix (and/or outer fixes or arcs based on that meter fix) and at the runway.

Scheduling is done according to a modified first-come-first-served order. Aircraft are sequenced according to their estimated time of arrival (ETA) at the meter fix, with modifications to this order made for manual sequence constraints entered by the traffic managers. An initial schedule is then created that conforms to acceptance rate constraints and to FAA aircraft spacing requirements.

This schedule is used to predict arrival times at the runway. These times are deconflicted according to wake vortex requirements, acceptance rate constraints, and runway occupancy restrictions. Any delay in excess of that which the TRACON can absorb is passed back, delaying the STAs at the meter fix. This additional delay will then have to be absorbed by the ARTCC feeding that meter fix.

TMA-SC infrastructure limitations

TMA-SC, however, has several important limitations. First, TMA-SC was not built to accept more than a single source of data input (the ARTCC's Host Computer System – HCS), although the FAA has recently begun to investigate providing data from two ARTCCs to a TMA-SC through the "Adjacent Center Data Feed" project.

This minimal connectivity between ARTCCs is a significant limitation. Airports in the Northeast corridor often lie close to the border of two or more ARTCCs. For example, Philadelphia International Airport (PHL), the airport chosen as the target for concept tests, sits on the border of the New York ARTCC (ZNY) and the Washington ARTCC (ZDC), and is within 130 nautical miles (nm) of the Boston (ZBW) and Cleveland (ZOB) ARTCCs. All four of these ARTCCs are responsible for controlling traffic flows into PHL, and are directly impacted by traffic flow management initiatives for PHL.

In addition, many Northeast corridor sectors are small and handle complex traffic problems (heavy crossing flows of traffic, climbing/descending traffic, and coincident streams of traffic destined for different airports). Because of these complications, controllers of such sectors are not able to impart significant delay to any particular aircraft, since they do not have either the physical space to lengthen the path of the aircraft by vectoring, nor can they spare the workload to track an aircraft deviating significantly from its flight plan. Slowing an aircraft is also less effective since there is less path distance over which this speed change can affect an aircraft's ETA.

If the amount of delay that needs to be absorbed remains the same, but the amount of delay per sector decreases, then, more sectors are needed to absorb the delay. Practically, this means spreading the delay further upstream from the airport, most likely across multiple ARTCCs as well. This requirement is also relevant to other areas within the National Airspace System (NAS) where earlier operational assessments found that delays often propagate upstream of a congested resource up to 500-1,000 nm away and across multiple ARTCCs.[4]

The proximity of multiple ARTCCs and the need to distribute delay across many sectors means that several ARTCCs must cooperate in meeting the STAs generated by a time-based metering system in the Northeast corridor. Moreover, the system that enables time-based metering must utilize data from these ARTCCs.

Not only must a time-based metering system for the Northeast corridor accept data from multiple HCSs, but in addition, it needs to accept data from another system, called the Enhanced Traffic Management System (ETMS). This is because a great deal of traffic bound for airports in the

Northeast corridor never enters ARTCC airspace. These "tower enroute control" aircraft depart from one airport and land at their destination by traversing adjacent TRACONs. The flight plans and radar track information for these aircraft are not stored in any HCS; rather, the data resides in the TRACON facility and is available from the ETMS computer. A Northeast corridor time-based metering system must be able to access this information.

Freeze horizon

The requirement that delay be distributed further upstream conflicts with the desire to minimize the distance over which the system needs to accurately predict trajectories. For the TMA-SC implementation, this trajectory prediction horizon is defined by the "freeze horizon".

The freeze horizon is the distance or time from the meter point at which an aircraft's schedule stops being updated. When the aircraft closes to within the specified distance from the meter point, the scheduled time is "frozen", and the scheduler no longer updates that aircraft's STA to that meter point. The STA would also be frozen if the aircraft is not tracked by a specified time prior to its STA. This situation would occur for an aircraft at an airport close to the meter point that has not yet departed.

Freeze horizons are critical for assuring that the STAs displayed to controllers will not change as they are controlling an aircraft. Without "freezing" the aircraft's STA, the controller would have a moving target as an STA.

Freezing the STA, however, means that any inaccuracies in the ETA to its meter point are also frozen. If the ETA is poor at the time the aircraft's STA is frozen, then the STA may also be poor. Poor STAs can result in improper sequences of aircraft, or difficulty in meeting STAs. For these reasons, it is desirable to keep freeze horizons as short as practical.

For TMA-SC, the freeze horizon is about 200 nm from the airport being metered. For the Northeast corridor, it is expected that aircraft will need to be metered over a distance of about 400 nm to achieve the necessary delay absorption. It was found that trajectory prediction was not sufficiently accurate over this greater distance, which suggested that the scheduling functions needed to be distributed into smaller segments.

Initial NASA Ames simulations

This issue was tested through controller-in-the-loop simulations conducted at NASA's Ames Research Center from October 2000 through June 2001. These simulations began with the concept of simply applying TMA-SC to the northeast (and to Philadelphia International Airport in particular), but it quickly became apparent that the airspace's limitations made this approach untenable.

The initial approach attempted to combine the four ARTCCs (ZNY, ZDC, ZOB, and ZBW) into one "super center" with a 400 nm freeze horizon. As with TMA-SC, a schedule was computed to the runway and the meter fixes at the boundary of the TRACON. Controller participants in the simulations were shown STAs for two points: the meter fixes, and the border of ZNY with both ZOB and ZBW. This latter time was computed by offsetting the meter fix time by the flying time between the ZNY border and the meter fixes.

The controllers then provided instructions for "pilots" (actually pseudo pilots controlling a number of aircraft in the simulation) to meet these STAs. Delay could therefore be attributed to two places: inside ZNY (by meeting a time at the meter fix), or outside of ZNY (by meeting a time at the border of ZNY with either ZOB or ZBW. Since ZNY controllers indicated that they could only absorb one or two minutes of delay inside their center, any delay in excess of that amount would have to be taken by the controllers in ZOB and ZBW.

Controller participants from ZOB sectors near the ZNY border reported that the sector sizes and the nature of the traffic flow through their sectors would not allow them to absorb the amount of delay that the system was suggesting. One result was therefore that it would be necessary to pass the delay farther upstream than the sectors close to the border with ZNY.

A second result from these initial simulations was that freezing the STAs 400 nm from the TRACON resulted in incorrect sequencing and poor STAs. The ETAs at 400 nm were inaccurate, which resulted in inaccurate STAs. While the aircraft traversed the 400 nm across the system, the ETAs would update and become more precise, but since the aircraft were within the freeze horizon, the STAs would not reflect these more accurate ETAs. The result of these poor ETAs was that the aircraft sequence shown by the software did not match the sequence expected by the controllers. The delay values for these out-of-sequence aircraft were thus either illogical or unattainable.

Subsequent simulation models attempted to break the problem up into two pieces. Instead of one large system, where STAs are frozen 400 nm from the airport based on ETAs to the runway, two TMA-SC systems were run. One system (the "inner" system) was configured to generate STAs to the runway and meter fixes at the PHL TRACON boundary. This system had a freeze horizon of 200 nm. A second system (the "outer" system) was configured to produce STAs at the runway and at metering points close to the boundaries between the ARTCCs (but not at the boundary of the PHL TRACON). The freeze horizon for the outer system was 400 nm from the runway, but only 200 nm from the meter points. Controllers in ZOB and ZBW delayed aircraft to meet the STAs to the ZNY border (generated by the outer system). They then handed off the aircraft to the ZNY controllers who delayed aircraft to meet times at the meter fixes (generated by the inner system).

This configuration resulted in better STAs and better sequences, but the two schedules were not communicating, resulting in inconsistencies. The improvement in performance, however, suggested that a distributed, networked scheduling architecture would produce better results.

This need for a distributed-scheduling architecture is consistent with the findings of earlier assessments of en route traffic operations outside of the Northeast corridor. The general need for distributed scheduling was formulated to support en route metering on a regional basis (between and across ARTCCs).[5] The concept of distributed scheduling is specifically adapted here for the implementation of arrival metering within the Northeast corridor.

ARCHITECTURE

Since the completion of the initial round of simulations, the distributed scheduler concept has been refined and implemented into the TMA-MC system. The following is a description of the basic system architecture.

Requirements

The architecture of a time-based metering scheduler for the Northeast corridor must conform to a number of requirements:

- the system should be compatible with TMA-SC to the extent possible in order to make use of proven and developed technology,
- the system should not require a great deal of interprocess communication,
- the system should provide ETA accuracy of about plus or minus 1 minute across its freeze horizon,
- the system should provide accurate sequences of aircraft,
- the system should distribute delay in accordance with sector delay absorption capacity, and
- the system should, as much as possible, not change the basic roles and responsibilities of the ARTCC controller for a time-based metering environment.

Frequently, scheduling systems attempt to define an "optimal" solution. Such a solution would result in a string of tightly coupled STAs. Such a solution has two major problems: it would require a great deal of interprocess communication, and ETAs and sequences would need to be accurate over 400 nm.

In contrast, the goal of distributed scheduling is to set up a loosely coupled schedule at each meter point. In doing so, controllers in these sectors, by meeting the computed STAs, would pass along a manageable, but not fully solved, problem to the downstream sector. Such a system would require less interprocess communication, and could be broken into pieces such that ETAs and sequences need only be accurate over shorter distances.

For the purposes of distributed scheduling, manageable is defined as not exceeding the "allowed maximum delay time" (AMDT) of the sector. AMDT is a predetermined maximum amount of delay time that a sector controller can absorb through vectoring and/or speed changes of an aircraft. It will vary according to the density of traffic in a sector, the complexity of the traffic flows through the sector, the physical amount of airspace available, and other factors. This value is determined by the traffic managers in the facility in which the sector resides in consultation with the sector controllers.

The design chosen is a network of schedulers that pass back restrictions in the form of rates. One scheduler would pass back a maximum number of aircraft that it could receive from an upstream scheduler within a period of time.

It is important to note that these rates would not be associated with particular aircraft; upstream schedulers would be free to choose which aircraft end up satisfying the rate. This allows sequence changes and schedule fluctuations to occur without affecting downstream schedules.

System concept

To explain the design, a generic airport arrival route is shown in Figure 1, consisting of a runway and a number of meter points (D1, C1, C2, B1, and A1). A similar sequence of points would occur upstream of points A2 through A4, but, for the purposes of the example, only the A1 branch will be examined.

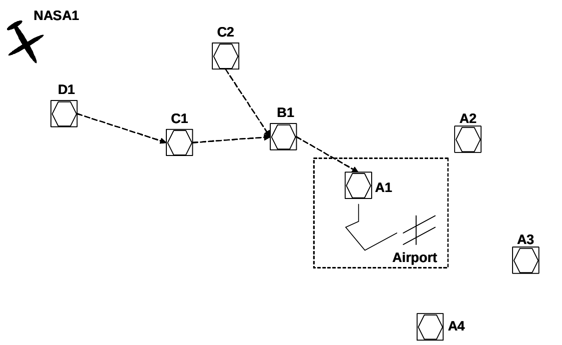

Figure 1. Generic metering topology.

Aircraft traverse through the set of meter points to get to the runway. These points are generally in different sectors and/or ARTCCs, and are placed at sector boundaries or merge points. Sector boundaries are efficient since they give the controller in that sector the maximum amount of airspace in which to delay an aircraft. Merge points are also efficient because the controller must join the streams of aircraft at these points, and providing times that are deconflicted reduces controller workload.

Consider an aircraft (NASA1) going from D1 to C1 to B1 to A1 to the runway. If sufficient demand is placed on any point along its route of flight, then NASA1 will need to incur delay. Typically, the most congestion occurs at the runway.

In order to predict demand and come up with a schedule, ETAs for NASA1 are calculated for the meter points along its route of flight, including D1, C1, B1, A1, and the runway. Similarly, ETAs are calculated for all the arrival aircraft flying in to these points.

The tightly coupled approach is to take these times and deconflict them centrally. However, the complexity of coordinating schedules, as well as the need to communicate so many ETAs across different ARTCCs, makes this approach impractical. In addition, such a schedule would be highly dependent on ETAs to the source of the majority of the congestion – the runway. Due to routing, speed, and altitude uncertainty, estimates of more distant points (such as the runway) are less precise than for closer points (such as D1). This means that the runway is the point to which the system has the least precise ETA.

A more loosely coupled approach is to break the problem into a series of smaller scheduling problems, and create constraints at each point that must be complied with by upstream schedulers feeding that point. This approach is depicted in Figure 2.

In Figure 2, the runway scheduler would calculate restrictions with which A1, B1, C1, C2, and D1 should comply in scheduling aircraft; A1 would calculate similar constraints for points upstream of it, and so on. The constraints can be designed to be generic (i.e. not attached to particular aircraft) so that particular sequences need not be adhered to by the controller in meeting these constraints. This makes the system more robust to poor ETAs by allowing sequences to change between meter points.

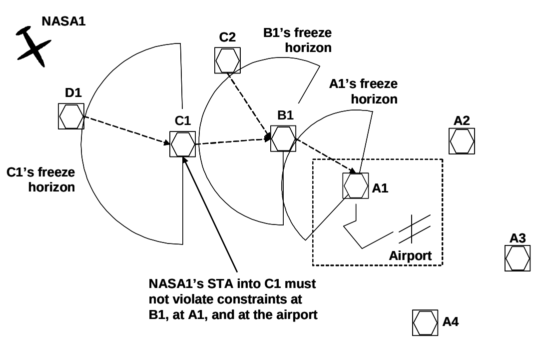

Figure 2. Distributed scheduling topology.

Rate profiles

The constraints that are passed between schedulers are called "rate profiles". Each rate profile contains a sequence of rates. These rates identify the number of aircraft that can be accepted from a particular stream of aircraft in a given period of time. It is calculated by counting the number of scheduled aircraft that originated from a particular place upstream which have been scheduled at that point. An example is shown in Figure 3 for B1's rate profile for point C1 and C2.

Figure 3 shows a timeline from 15:00 to 15:15 for the point B1 (Figures 1 and 2) that is fed by points C1 and C2. The timeline is read from the bottom, with the earliest time (15:00) at the bottom of the timeline, and the latest time (15:15) at the top of the timeline. The arrows indicate the STAs for various aircraft on the timeline, and the labels to the right of the vertical timeline indicate the source of the aircraft (either C1or C2).

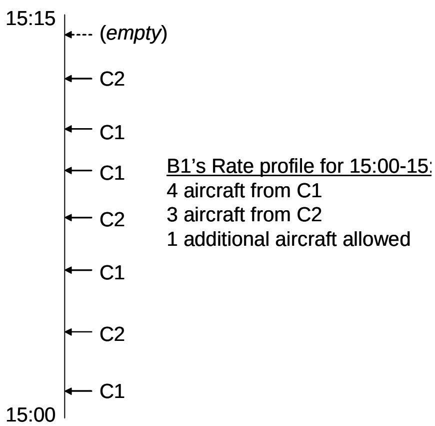

Figure 3. Rate profile example.

Since these aircraft have been successfully scheduled, the count shown in Figure 3 (4 aircraft from C1 and 3 from C2) represents the minimum number of aircraft from C1 and C2 respectively that can arrive at B1 within the given time period. Any unused capacity of this time period can also be determined and identified to all upstream points as being available. In this way the total capacity of a meter point is divided between streams entering the fix.

Such a time period is called a "bin". Bins can be labeled using the hour and an index, so that the bin from 15:00 to 15:15 would be called bin 15-01.

Calculation of additional capacity

While the count of aircraft represents the minimum, any "holes" in the bin (slots in which aircraft can be, but have not been, scheduled) must be accounted for since they may need to be used by new aircraft entering the system or by aircraft being rescheduled into a new bin.

For example, suppose we have the situation shown in Figure 4. We have 6 aircraft (NASA1 through NASA6) scheduled, resulting in the gaps shown. We check the size of the gaps for the ability to fit in an aircraft given a MIT restriction at the fix (assume it is 6 MIT for this example). If we find a gap between two aircraft in which another aircraft could fit, we know we have more capacity in that bin than demand.

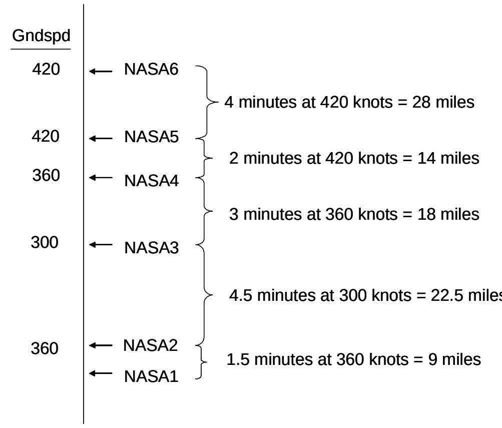

Figure 4. Rate profile hole counting.

Going farther, if provided with the aircraft speeds and MIT restrictions, we could count those "holes": with a 6 MIT restriction no aircraft can be scheduled between NASA1 and NASA2, two can be scheduled between NASA2 and NASA3, two between NASA3 and NASA4, one between NASA4 and NASA5, and three between NASA5 and NASA6. In general, the number of "holes" between two aircraft is:

$$\frac{Separation\_in\_miles}{MIT\_restriction} - 1 \qquad [1]$$

Note that the result of Equation 1 may not be a whole number, and requires some accounting through a fractional-hole-cleanup process. For example, a more precise accounting of the holes shown in Figure 4 would result in the totals shown in Table 1.

Table 1. Hole accounting for Figure 4.

Hole	Calculation	Result
NASA1 to NASA2	9/6 - 1	0.5
NASA2 to NASA3	22.5/6 - 1	2.75
NASA3 to NASA4	18/6 - 1	2.0
NASA4 to NASA5	14/6 - 1	1.33
NASA5 to NASA6	28/6 - 1	3.66
Total		10.25

In the previous example (Figure 4), it was assumed that NASA1's STA marked the start of the bin, and NASA6's STA marked the end of the bin.

However, this will usually not be the case. There will typically be holes prior to the first aircraft in the binand after the last aircraft in the bin. Interpolation can be used for accounting for these gaps.

For example, Figure 5 depicts a gap between NASA1 and NASA0 that straddles two bins. The gap between the aircraft has a start and end time, as does each bin. Calculation of how much of this gap is attributable to one bin or the other is accomplished using Equation 2, filling in the values for the bin for which the accounting is desired.

$$bin\_factor = \frac{\min(gap\_end, bin\_end) - \max(gap\_start, bin\_start)}{gap\_end - gap\_start} \qquad [2]$$

This factor is then multiplied by the size of the gap between the aircraft (as given by Equation 1) to derive the interpolated number of holes in each bin.

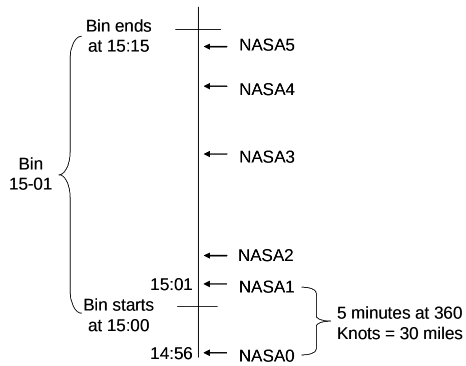

$$bin\_factor = \frac{15:01 - 15:00}{15:01 - 14:57} = \frac{1}{4} = 0.25$$

$$holes\_for\_bin\_15-01 = 0.25 \times \left(\frac{30}{6} - 1\right) = 1$$

Figure 5. Bin count interpolation example.

Scheduling an aircraft into a hole

Once an aircraft (such as “NASA7” in Figure 6) is actually scheduled into a hole, its scheduled time may affect other aircraft on the timeline. For example, in Figure 6, NASA3 would have to be delayed an additional 12 seconds to achieve the required 6 MIT separation. If NASA3’s STA were already frozen, its STA could not be changed, so NASA7 would have to be scheduled to arrive after NASA3.

Use of the rate profiles

In practice, if the scheduler at C1 (Figures 1 and 2) obtains the rate profile described in Figure 3, it will ensure that no more than 5 (4 scheduled plus the 1 additional) aircraft being scheduled into C1 will arrive at B1 between 1500 and 1515. Likewise, C2 would ensure that it scheduled no more than 4 (3 scheduled plus the 1 additional). The traffic arriving from C1 and C2 will then satisfy the constraint at B1.

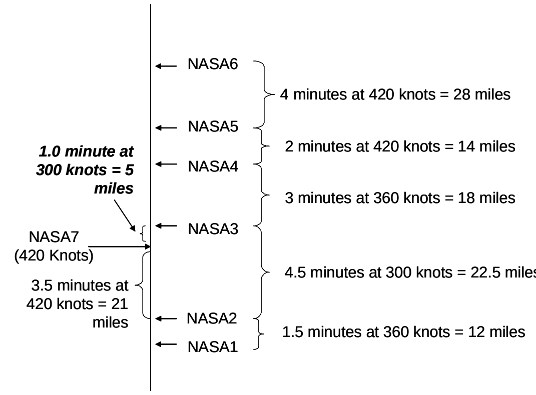

Figure 6. Inserting a new aircraft example.

It is possible that both C1 and C2 could attempt to use the same hole for scheduling an aircraft (the 1 additional slot at B1). In such a case, the schedulers at C1 and C2 would schedule the aircraft assuming they could use the additional slot. The scheduler at B1 would choose one aircraft or the other (based on a first-come, first-served rule). The scheduler at B1 would then update its rate profile, increasing its allocation of slots to either C1 or C2, and reducing the number of additional slots by 1 (to zero). During the next scheduling cycle one of the schedulers at C1 and C2 would see an increase in allocated slots, and both would see zero additional slots. Whichever scheduler was not allocated the additional slot would then have to delay an aircraft so as not to arrive in the now full bin. Any overlap in use of the same hole would therefore be corrected in one scheduling cycle (which should be within 10-20 seconds).

Distribution of the rate profiles

It is not sufficient for a meter point scheduler to check only the meter point immediately downstream of it. Each point must check the rate profiles for all meter points downstream of it. So, for example, the scheduler at C1 must check rate profiles at B1, A1, and the runway.

This requirement is driven by the fact that capacity is not strictly distributed upstream. The capacity of a bin at B1 is not distributed to specific bins at C1 and C2. Therefore, the successful scheduling of an aircraft in a bin at C1 does not guarantee successful scheduling at B1.

For example, consider that NASA1 is being scheduled into D1. The scheduler at D1 proposes an STA at D1, and then uses the flying time from D1 to C1 to get a proposed arrival time at C1. This arrival time at C1 is in bin 15-01. The scheduler at D1 checks the rate profile for C1 and determines that there is sufficient capacity to schedule the aircraft to arrive in bin 15-01. However, NASA1's arrival at B1 may put it in a bin that has no capacity – there is no way to know except by directly checking the rate profile at B1. The same is true for the rate profiles at A1 and the runway.

Coupling between schedules

For a tightly coupled schedule, aircraft would be deconflicted at all points. The most restrictive time would drive the scheduled time of arrival at all upstream points. This point would most likely be the runway, since that is where all the aircraft will ultimately merge. However, that ETA is the most distant, and most likely the least precise. Such a schedule would therefore be based upon the least precise ETA for that aircraft. This could cause poor sequencing, unnecessary delays, or too little delay (which could saturate the TRACON).

For a loosely coupled, rate profile-based system, ETAs beyond the next meter point are used only to check that an arrival rate restriction is not being violated. Poor ETAs could potentially cause an aircraft to be initially scheduled in an incorrect bin.

However, since particular aircraft are not assigned to particular downstream bins, many of these errors will not affect upstream STAs. This is particularly true since ETA error at any given point is expected to be a random, zero-mean process. With such a process, STA errors will balance out, with some errors causing decreases in rate profiles, others causing increases. Since rate profiles only reflect a sum, these errors cancel out.

For those that do not cancel out, the rate profile approach accommodates larger errors. For the tightly coupled approach, downstream STA errors in excess of the expected performance in meeting STAs (most likely plus or minus a minute) are directly reflected in upstream STAs. For the loosely coupled, rate profile approach, downstream STA errors must cause changes to rate profiles. Such changes require errors of up to the size of the bin (expected to be about 15 minutes). Therefore, larger errors in downstream ETAs are tolerated.

Rolling freeze

The loosely coupled architecture and multiple freeze horizons result in a "rolling freeze" effect, whereby an aircraft is frozen with respect to each meter point separately, as shown in Figure 7.

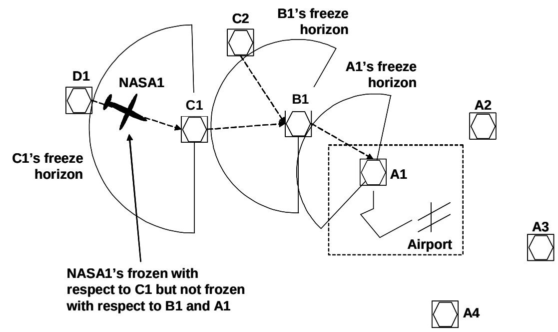

Figure 7. Rolling freeze example.

In Figure 7, NASA1's schedule at C1 is no longer being updated since it is inside the freeze horizon. This means that any corrections to the ETA at C1 will not result in schedule changes. However, since NASA1 is outside of B1's and A1's freeze horizons, its STAs at B1 and A1 will continue to be updated. ETA corrections to those points will continue to be reflected in the STAs until NASA1 freezes with respect to those points. This results in shorter distances over which ETAs must be highly accurate. It also means that the sequences of aircraft at downstream meter points can continue to change while ETAs are being updated.

IMPLEMENTATION

TMA-MC is currently undergoing testing at NASA Ames Research Center, at Computer Science Corporation's offices in New Jersey, at the FAA's William J. Hughes Technical Center, and at field sites within the Boston, Cleveland, New York, and Washington ARTCCs. When deployed, TMA-MC will enable coordinated time-based metering to airports whose proximate arrival routes traverse multiple ARTCCs. The generic architecture described in earlier sections has been implemented in the TMA-MC system and adapted to the Northeast corridor.

TMA-MC hardware

The TMA-MC system is a networked system of UNIX workstations.[6] Each participating ARTCC has a suite of equipment connected on a Wide Area Network. The TMA-MC equipment communicates with the ARTCC's HOST computer over the HOST Interface Device/National Airspace System Local Area Network. This connection allows the STAs generated by the scheduler to be displayed on the controllers' displays.

To manage time-based metering operations, users from the Traffic Management Unit (TMU) are given two main interfaces within the Traffic Manager Graphical User Interface (TGUI). The TGUI provides Traffic Management Coordinators with a timeline display of the schedules into meter points within their ARTCC, and with a display of the timelines to the runway and meter points on the boundary of the TRACON. The TGUI also provides a Load Graph display that shows a provides a graphical depiction of the future traffic demand and delay at meter points chosen by the user.

TMA-MC software

The TMA-MC software has been adapted from the TMA-SC software, and is backward compatible with the TMA-SC architecture. Since TMA-SC is continually being upgraded, routine resynchronizations with TMA-SC are conducted to ensure continued compatibility and to ensure that new TMA-SC features are incorporated into TMA-MC. This allows for easy transition from TMA-SC to TMA-MC if desired.

The core processes for both the TMA-MC and the TMA-SC system can be grouped into: user interfaces, communication processes, processes designed to determine aircraft trajectories and ETAs, and the dynamic planner. These processes are described in detail in other publications.[7] For the remainder of this paper, the trajectory and ETA calculation processes will be grouped with the dynamic planner as part of a general scheduling process.

TRACON dynamic planner

The TRACON dynamic planner (DP) calculates and provides STAs to the runway and to the meter points at the boundary of the TRACON. It is essentially the same as that used by TMA-SC. Additional software modules have been added in order to calculate rate profiles and handle the additional communications required by TMA-MC.

Single point dynamic planner

The single point dynamic planner (SPDP) is an adaptation of the original TMA-SC DP. The SPDP utilizes the scheduling algorithms described in this document, enabling local scheduling to a meter point while respecting rate profiles of downstream meter points. It also calculates and distributes rate profiles for its meter points.

Topology of TMA-MC airspace for PHL

The TMA-MC system for PHL utilizes a hierarchical topology, as shown in Figure 8. At the center of the system, handling the scheduling for the PHL TRACON, is a TRACON DP. This scheduler provides STAs to the runway and to the meter fixes labeled as SPUDS, PTW, BUNTS, TERRI, and VCN.

Upstream of the TRACON are meter points. These locations are the points at which SPDPs will deconflict and schedule aircraft. These points can be considered to be in tiers beyond the TRACON meter points. First tier points in Figure 8 are labeled as DNY, HAR, OTT, ORF, and DUNEE. ZOB has two additional tiers, a second tier consisting of the point labeled JST, and a third consisting of the point labeled DJB.

In many cases it is desirable to have STAs at points other than the meter points, such as at the points indicated by the narrow gray rectangles in Figure 8. These locations are referred to as “outer windows”. In this case scheduled times can be obtained by subtracting the time-to-fly from the desired point to the next meter point.

Sample scheduling scenario

As an example, consider the metering configuration shown in Figure 8. The system is scheduling aircraft and providing STAs to the air traffic controllers. An aircraft, NASA1, is flying from DJB (a fix named Dryer) to JST (Johnstown) to HAR (Harrisburg) to BUNTS to PHL. There are DPs for each meter point (SPDPs for JST, HAR, and the DP for the TRACON).

Each time a track hit from an ARTCC’s radar is recorded, NASA1’s position, altitude, and speed are updated. This causes the system to update NASA1’s ETAs to JST and the border of ZOB’s airspace with ZNY. ZNY’s system then takes the updated border crossing time and updates NASA1’s ETA to HAR, BUNTS, and PHL.

The DP scheduling JST takes the new ETA for NASA1 to JST and deconflicts it locally (with other aircraft over JST). It has a local copy of the rate profile that HAR has allocated for JST aircraft, a copy of the rate profile that BUNTS has allocated to JST, and a copy of the rate profile that the runway has allocated to JST.

The JST DP then estimates NASA1’s flying time to JST, to BUNTS, and to the runway, and checks

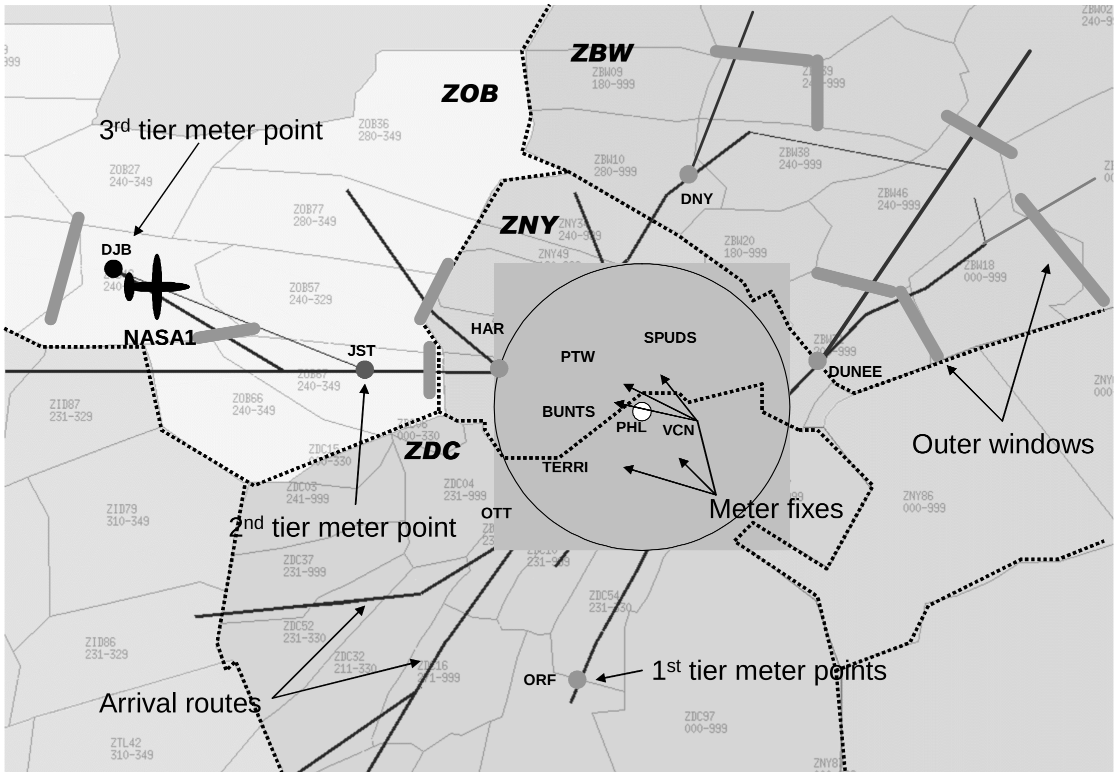

Figure 8. Topology of TMA-MC airspace for PHL.

to see if there is room in the bin in which NASA1 will arrive at each of those points. If there is no room, then the amount of delay needed to put the aircraft in to the next available bin will be determined.

That delay is apportioned to the points between NASA1's position and the meter point at which the delay needs to occur, starting with the meter point furthest downstream. The delay is allocated so as to not exceed the AMDT of the sectors between the meter points. If the needed delay exceeds the sum of the AMDTs of all the sectors, the additional delay is added to the point being scheduled (in this case JST).

If delay has been added to the current meter point, then this new scheduled time must then go through the rate profile checking process again. Iterations of this checking process stop when a time has been found that does not require additional changes to the STA.

<u>NASA1's progress through the system</u>

Figure 8 depicts NASA1 in a position such that its STA would be frozen with respect to JST, but not with respect to HAR or BUNTS. This means the controller for NASA1's current sector has a time for NASA1 to meet at the next meter window (which is near the boundary of that ARTCC controller's sector with the adjacent ARTCC sector). If the aircraft meets that STA, then NASA1 can be easily merged with the flows into JST, and will not need to be delayed beyond the AMDTs of any downstream sectors.

The freeze horizon for HAR is set to be just beyond the next JST meter window for NASA1. Prior to that point, NASA1's STA at HAR will continue to be updated. Once past that point, NASA1's STA with respect to HAR will be frozen. The controller in the next sector will try to have NASA1 meet its STA at the HAR meter window between HAR and JST. This window is set at the eastern border of that sector. This configuration allows the controller for that sector the entire length of his or her sector to delay NASA1 (if necessary) through vectoring or speed control.

The HAR meter window is also set at the border of ZOB with ZNY. Before the aircraft enters the airspace of ZNY, control of the aircraft will be "handed off" to ZNY controllers. Once handed off, NASA1 will then be vectored or slowed by ZNY controllers to meet an STA at BUNTS.

FUTURE WORK

Simulation and field test activities

The algorithms described in this document are undergoing testing at the NASA Ames Research Center as part of the TMA-MC system. A simulation capability for closed-loop testing of TMA-MC has recently been developed and testing of the scheduling algorithm has commenced. The results of these tests will be used to evaluate the accuracy, efficiency, and robustness of the algorithm under varying levels of STA compliance.

Following these closed loop tests, controller-in-the-loop testing will begin. These tests, scheduled to start in January 2004, will provide training to controllers in the use of the system. In addition, these tests will examine the overall usability of the system, the ability of the specific sectors to absorb delay, and the coordination required between ARTCCs and sectors within the ARTCCs.

Field trials of TMA-MC are scheduled for spring, 2004. These field trials will begin with shadow testing ofthe system in the four Northeast corridor ARTCC facilities. For these shadow tests, ARTCC personnel will monitor the system's displays before making any operational decisions based on it. When confidence is gained that the system will operate effectively, progressively more use of the system will be made. The goal of these field trials is to actively meter live traffic using TMA-MC advisories. It is hoped that data derived from these field trials will demonstrate the benefits of the TMA-MC system.

Future developments

Once TMA-MC has been satisfactorily tested, NASA will turn over the technology to the FAA, who decides whether to deploy the system. Although TMA-MC is being developed for testing at PHL, the flexibility of the system allows it, with some modifications, to be used for other airports that lie close to ARTCC boundaries.

Enhancements to TMA-MC are planned to extend time-based metering throughout the NAS for managing en route operations on a regional basis. These enhancements will add the capability to fill three key gaps in the support of regional traffic flow management:

- En route metering of arrivals upstream of the extended terminal area served by TMA-SC/MC (typically 200-400 nm or 20-30 minutes upstream of the airport)
- Arrival metering to small terminals that do not warrant a full terminal-metering adaptation
- En route metering of traffic to congested airspace independent of destination airport

These enhancements will enable each ARTCC to simultaneously manage multiple restrictions to one or more downstream facilities. A distributed network of TMA capabilities across ARTCCs will allow traffic flow management specialists to "daisy chain" their capabilities to manage restrictions that propagate upstream. This flexibility will enable traffic flow restrictions to be dynamically adapted to mitigate congestion with minimal impact on airspace users.

REFERENCES

[1] Free Flight Program Office, "Free flight program performance metrics results to date: December 2002 report," Federal Aviation Administration, Washington, DC, December 2002.

[2] Farley, T., J. Foster, T. Hoang, and K. Lee, "A time-based approach to metering arrival traffic to Philadelphia," First AIAA Aircraft Technology, Integration, and Operations Forum, Los Angeles, California, October 2001.

[3] Wong, G. L., "The dynamic planner: The sequencer, scheduler, and runway allocator for air traffic control automation" NASA/TM-2000-209586, April 2000.

[4] Green, S. M., and M. P. Grace, "Conflict-free planning for en route spacing: A concept for integrating conflict probe and miles-in-trail," AIAA 99-3988, AIAA Guidance, Navigation, and Control Conference, Portland, OR, August 1999.

[5] Titan Systems Corporation, "Regional metering operational concept description," Year 2002 Update, NRA RTO-72 CDRL:#3a.1, Advanced Air Transportation Technology (AATT) Office, NASA Ames Research Center, Mountain View, CA, 94035, October 2002.

[6] Hoang, T., T. Farley, J. Foster, and T. Davis, "The multi-center TMA system architecture and its impact on inter-facility collaboration," 2nd AIAA Aircraft Technology, Integration, and Operations Forum, Los Angeles, California, October 2002.

[7] Erzberger, H., Davis, T. J., and Green, S. M., "Design of center-TRACON automation system," Proceedings of the AGARD Guidance and Control Panel 56th Symposium on Machine Intelligence in Air Traffic Management, Berlin, Germany, 1993.